

On the boundary of the region defined by homomorphism densities

Hamed Hatami ^{*} Sergey Norin [†]

January 2, 2017

Abstract

The Kruskal–Katona theorem together with a theorem of Razborov [Raz08] determine the closure of the set of points defined by the homomorphism density of the edge and the triangle in finite graphs. The boundary of this region is a countable union of algebraic curves, and in particular, it is almost everywhere differentiable. One can more generally consider the region defined by the homomorphism densities of a list of given graphs, and ask whether the boundary is as well-behaved as in the case of the triangle and the edge. Towards answering this question in the negative, we construct examples which show that the restrictions of the boundary to certain hyperplanes can have nowhere differentiable parts.

1 Introduction

For two simple finite graphs F and G , the *induced subgraph density* of F in G , denoted by $p(F; G)$, is the number of induced copies of F in G divided by $\binom{|V(G)|}{|V(F)|}$. A large part of extremal graph theory is concerned with understanding the asymptotic relation between induced subgraph densities of different graphs. In other words, given a list of graphs $\mathcal{F} = \{F_1, \dots, F_k\}$, one would like to determine the compact region $T(\mathcal{F}) \subseteq [0, 1]^k$ defined as the set of limit points of sequences $\{(p(F_1; G_i), \dots, p(F_k; G_i))\}_{i=1}^\infty$, where $|V(G_i)| \rightarrow \infty$. The general problem of understanding the regions defined by subgraph densities was essentially first considered in [ELS79]. More recently, the study of graph limits [Lov12] has brought new attention to such problems.

Describing the regions $T(\mathcal{F})$, even when \mathcal{F} consists of a few small graphs, is a major challenge. To the best of our knowledge the only non-trivial cases for which a complete description is known are the following. The set $T(K_2, K_3)$ has only been completely described less than a decade ago, and Razborov’s solution [Raz08] to this problem is highly non-trivial, and is considered to be one of the major achievements of the theory of flag algebras [Raz07]. More recently, Razborov’s result was extended in two directions. Reiher [Rei16] determined the inequalities defining the set $T(K_2, K_t)$ for every t , while Glebov et al. [GGH⁺16] and Huang et. al [HLN⁺14] obtained the description of $T(\mathcal{F})$ for any pair \mathcal{F} of graphs on three vertices.

The structure of $T(\mathcal{F})$ is very complex in full generality. In particular, it is shown in [HN11] that the separation problem for $T(\mathcal{F})$ is undecidable. That is there exists no algorithm which given \mathcal{F} and an affine halfspace S of \mathbb{R}^k decides whether $T(\mathcal{F}) \cap S = \emptyset$ or not.

The focus of this paper is the boundary of the region $T(\mathcal{F})$. The boundary of $T(K_2, K_3)$ is a countable union of algebraic curves [Raz08], and, in particular, is almost everywhere differentiable. This raises the question that whether the boundary is always as well-behaved; for example, is it almost everywhere differentiable?

Towards giving a strong negative answers to such questions we give an example which shows that the intersection of $T(\mathcal{F})$ with a hyperplane can have a nowhere differentiable boundary. The precise statement of our main result, Theorem 4.1, is technical and requires some preparation. In Section 2 we introduce

^{*}School of Computer Science, McGill University. hatami@cs.mcgill.ca. Supported by an NSERC grant.

[†]Department of Mathematics and Statistics, McGill University. snorin@math.mcgill.ca. Supported by an NSERC grant.

the required terminology from the theory of graph limits. In Section 3 we prove some technical results, and finally in Section 4 we state and prove our main result. We finish by some concluding remarks in Section 5.

2 Preliminaries

In this section we introduce the notions from the theory of graph limits that we will be using for the rest of the paper.

All graphs considered in this paper are simple and finite. For a graph G , let $V(G)$ and $E(G)$, respectively denote the set of the vertices and the edges of G . The unique graph with no vertices is denoted by \emptyset .

2.1 Weighted graphs and graphons

We redefine $T(\mathcal{F})$ in terms of the induced homomorphism densities as follows. For two graphs F and G , a map $\phi : V(F) \rightarrow V(G)$ is a *strong homomorphism from F to G* if $vw \in E(F)$ if and only if $\phi(v)\phi(w) \in E(G)$; that is ϕ preserves both adjacency and non-adjacency. Let $s(F, G)$ denote the number of strong homomorphisms from F to G , and let *the induced homomorphism density*

$$t(F; G) := \frac{s(F, G)}{|V(G)|^{|V(F)|}},$$

be the probability that a random mapping (not necessarily injective) from the vertices of F to the vertices of G is a strong homomorphism. We define $t(\emptyset; G) := 1$, for every graph G . For $\mathcal{F} = \{F_1, F_2, \dots, F_k\}$, the set $T(\mathcal{F})$ defined in the introduction is (up to rescaling) equal to the closure of the set of points $(t(F_1; G), t(F_2; G), \dots, t(F_k; G))$ for all G .

A *weighted graph* (G, μ) is a graph G together with a discrete probability measure μ on $V(G)$, such that $\mu(v) > 0$ for every $v \in V(G)$. (The last assumption is unusual, but convenient.) The quantity $t(F; G, \mu)$ can be defined similarly to $t(F; G)$ by mapping each vertex of F to a vertex of G randomly and independently according to the distribution μ , rather than uniformly.

A sequence of graphs $\{G_i\}_{i=1}^\infty$ is called *convergent* if for every graph F , the sequence $\{t(F; G_i)\}_{i=1}^\infty$ converges. Weighted graphs are not sufficient to represent the limits of all convergent sequences. However, Lovász and Szegedy proved in [LS06] that the limit of a convergent graph sequence can be represented by a so-called *graphon*, which is a symmetric measurable function $W : [0, 1]^2 \rightarrow [0, 1]$. More precisely, given a convergent graph sequence $\{G_n\}_{n=1}^\infty$, there always exists a graphon W such that for every graph F , we have

$$\lim_{n \rightarrow \infty} t(F; G_n) = \mathbf{E} \left[\prod_{(u,v) \in E(F)} W(x_u, x_v) \prod_{(u,v) \notin E(F)} (1 - W(x_u, x_v)) \right], \quad (1)$$

where $\{x_u \mid u \in V(F)\}$ are independent random variables taking values uniformly in $[0, 1]$. We denote the expected value in the right-hand side by $t(F; W)$. Conversely, for every graphon W , one can construct a graph sequence that converges to W in the above sense. Note that a convergent sequence of graphs does not converge to a *unique* graphon. We say that graphons W and W' are *weakly isomorphic* if $t(F; W) = t(F; W')$ for every graph F . Thus a convergent sequence of graphs does converge to a unique graphon up to weak isomorphism. It will be convenient for us not to distinguish between weakly isomorphic graphons in this paper.

For every weighted graph (G, μ) (on vertex set $[n]$), we define a step-function $W_{G, \mu}$ as follows: split $[0, 1]$ into n intervals J_1, \dots, J_n of lengths $\lambda(J_i) = \mu(i)$, and for $x \in J_i$ and $y \in J_j$, set $W_{G, \mu}(x, y)$ to 1, if i is adjacent to j in G , and to 0, otherwise. Note that we have $t(F; G, \mu) = t(F; W_{G, \mu})$ for every graph F and weighted graph (G, μ) . In this sense graphons are a natural extension of weighted graphs (and, in particular, graphs).

2.2 Quantum graphs

Let L be a finite (possibly empty) subset of natural numbers. An L -labeled graph is a graph in which some of the vertices are labeled by distinct elements of L , such that every label is assigned to exactly one vertex (there may be any number of unlabeled vertices). Let \mathcal{H}_L denote the set of all L -labeled graphs up to label-preserving isomorphism. For brevity we denote the set of unlabeled graphs \mathcal{H}_\emptyset by \mathcal{H} . A graph is *partially labeled* if it is L -labeled for some set L . A partially labeled graph in which all vertices are labeled is called a *fully labeled* graph.

We extend the definition of the induced homomorphism density to partially labeled graphs in the following way. Consider a finite set $L \subset \mathbb{N}$, a partially labeled graph $H \in \mathcal{H}_L$, a graph G , and a map $\phi: L \rightarrow V(G)$. Then $t(H, \phi; G)$ is defined to be the probability that a random map from $V(H)$ to $V(G)$ is a strong homomorphism conditioned on the event that the labeled vertices are mapped according to ϕ .

The definition of $t(H, \phi; G)$ extends to graphons in a straightforward manner. That is, given a graphon W and a map $\phi: L \rightarrow [0, 1]$, the induced homomorphism density $t(H, \phi; W)$ is the expected value in (1) but now for every $i \in L$, we fix $x_{u_i} = \phi(i)$ where u_i is the vertex with label i , and the randomness is over the rest of the variables.

Consider real numbers $\alpha_1, \dots, \alpha_k$ and graphs H_1, \dots, H_k . In order to study the linear combinations of induced homomorphism densities, it is useful to define quantum graphs and labeled quantum graphs. An L -labeled quantum graph is an element of the vector space $\mathbb{R}[\mathcal{H}_L]$, i.e. it is a formal linear combination of graphs in \mathcal{F}_L . For a graph G , an L -labeled quantum graph, $f = \sum_{i=1}^n \alpha_i H_i \in \mathbb{R}[\mathcal{H}_L]$ and a map $\phi: \cup_{i=1}^k L_i \rightarrow V(G)$, define $t(f, \phi; G) := \sum_{i=1}^k \alpha_i t(H_i, \phi|_{L_i}; G)$.

We define the product $H_1 \cdot H_2$ of two vertex disjoint L -labeled graphs H_1 and H_2 , to be zero if the labeled subgraphs of H_1 and H_2 induced by the labeled sets differ, and, otherwise, we define $H_1 \cdot H_2$ to be the sum of all distinct L -labeled graphs obtained by identifying the vertices of H_1 and H_2 with the same labels and possibly adding some edges between unlabeled vertices of H_1 and unlabeled vertices of H_2 . This product is defined so that for every two partially labeled graphs $H_1, H_2 \in \mathcal{F}_L$, a graph G , and a map $\phi: L \rightarrow V(G)$,

$$t(H_1 \cdot H_2, \phi; G) = t(H_1, \phi; G)t(H_2, \phi; G). \quad (2)$$

We extend this product to $\mathbb{R}[\mathcal{F}_L]$ by linearity; that is if $f = \sum_{i=1}^n \alpha_i H_i \in \mathbb{R}[\mathcal{F}_L]$ and $g = \sum_{i=1}^m \beta_i F_i \in \mathbb{R}[\mathcal{F}_L]$, then $f \cdot g = \sum_{i=1}^n \sum_{j=1}^m \alpha_i \beta_j H_i \cdot F_j$. Consider a finite set $L \subset \mathbb{N}$, a graph G , and a map $\phi: L \rightarrow V(G)$. It follows from (2) that $f \mapsto t(f, \phi; G)$ defines a homomorphism from $\mathbb{R}[\mathcal{F}_L]$ to \mathbb{R} .

Let the linear map $\llbracket \cdot \rrbracket : \bigcup_L \mathbb{R}[\mathcal{F}_L] \rightarrow \mathbb{R}[\mathcal{F}]$ be defined by un-labeling all the labeled vertices. Note that $t(\llbracket f^2 \rrbracket, W) = 0$ for a graphon W and an L -labeled quantum graph $f \in \mathbb{R}[\mathcal{F}_L]$ if and only if $t(f, \phi; W) = 0$ for almost every $\phi: L \rightarrow [0, 1]$.

2.3 Infinite lexicographic products

The lexicographic product $F \otimes H$ of graphs F and H is the graph with vertex set $V(F) \times V(H)$ where two vertices (u, v) and (x, y) are adjacent in $F \otimes H$ if and only if either u is adjacent to x in F , or $u = x$ and v is adjacent to y in H . The lexicographic product of two weighted graphs is defined by

$$(F, \mu) \otimes (H, \nu) = (F \otimes H, \mu \times \nu).$$

Given an infinite sequence of weighted graphs $(F_i, \mu_i)_{i \in \mathbb{N}}$, the (infinite) lexicographic product $\otimes_{i=1}^\infty (F_i, \mu_i)$ is the graphon defined (up to weak isomorphism) as the limit of the convergent sequence of weighted graphs $\{\otimes_{i=1}^n (F_i, \mu_i)\}_{n \in \mathbb{N}}$. If all the graphs in the sequence $(F_i, \mu_i)_{i \in \mathbb{N}}$ are equal to the weighted graph (F, μ) then the resulting product is called the *infinite lexicographic power* of (F, μ) and is denoted by $\otimes^\infty (F, \mu)$.

Let W be a graphon, and let (F, μ) be a weighted graph. We say that a partition $(J_v)_{v \in V(F)}$ of the interval $[0, 1]$ is an (F, μ) -partition of W if

- $\lambda(J_v) = \mu(v)$ for all $v \in V(F)$, and
- for every pair of distinct vertices $u, v \in V(F)$, the restriction of W to $J_u \times J_v$ is identically 1, if u and v are adjacent in F , and is identically 0 otherwise.

Note that this implies no restrictions on the value of W on the diagonal cells $J_u \times J_u$ for $u \in V(G)$. We say that a graphon W is (F, μ) -*partitionable* if it admits an (F, μ) -partition.

For a graphon W and a measurable set $J \subseteq [0, 1]$ we define the *restriction* $W[J] : [0, 1]^2 \rightarrow [0, 1]$ of W to J by taking any measurable bijection $\phi : [0, 1] \rightarrow J$ such that $\lambda(S) = \lambda(\phi^{-1}(S))\lambda(J)$ for every measurable $S \subseteq J$, where λ denotes the Lebesgue measure, and defining $W[J](x, y) = W(\phi(x), \phi(y))$. The choice of ϕ is irrelevant for our purposes as the graphons obtained using different choices are weakly isomorphic to each other.

Given an (F, μ) -partition $(J_v)_{v \in V(F)}$ of a graphon W , let $W_v = W[J_v]$ for $v \in V(G)$. We say that $(W_v)_{v \in V(F)}$ is the collection of the *bags* of the (F, μ) -partition $(J_v)_{v \in V(F)}$.

It is important for the application of the above definitions to note that for any infinite lexicographic product $W = \otimes_{i=1}^{\infty} (F_i, \mu_i)$ of weighted graphs and every $n \geq 1$, there exists a $\otimes_{i=1}^n (F_i, \mu_i)$ -partition of W such that all the bags of the partition are weakly isomorphic to $\otimes_{i=n+1}^{\infty} (F_i, \mu_i)$.

2.4 Blowups

In addition to lexicographic products, in some of the technical arguments we will use another, simpler, construction, called blowups. For a graph F and a positive integer vector $\mathbf{a} \in \mathbb{N}^{V(F)}$, the \mathbf{a} -*blow-up* of F , denoted by $F^{(\mathbf{a})}$, is the graph obtained by replacing every vertex v of H with $\mathbf{a}(v)$ different vertices where a copy of u is adjacent to a copy of v in the blow-up graph if and only if u is adjacent to v in F . If \mathbf{a} is equal to d in every coordinate we will denote \mathbf{a} -blow-up of F by $F^{(d)}$.

If F is a fully labeled graph then the blow up $F^{(\mathbf{a})}$ is defined in a similar manner, where the newly added vertices will be unlabeled and the original vertices of F will keep their labels. Let $\tilde{F}^{(\mathbf{a})}$ be the quantum graph equal to the sum over all graphs that can be obtained from $F^{(\mathbf{a})}$ by possibly adding some edges between different copies of the same vertex of F .

Finally, for $I \subseteq V(F)$, let F^I denote $F^{(\mathbf{a}_I)}$ where $\mathbf{a}_I(v) = 2$ if $v \in I$, and $\mathbf{a}_I(v) = 1$, otherwise. That is, F^I is obtained from F by replacing every vertex in I by a pair of twin vertices. Similarly \tilde{F}^I is defined as $\tilde{F}^{(\mathbf{a}_I)}$.

2.5 Finite forcibility

A graphon W is *finitely forcible* if there exists a quantum graph f such that $t(f; W') = 0$ for a graphon W' if and only if W' is weakly isomorphic to W . Finitely forcible graphons, first formally defined in [LS11], are the structures that can appear as “unique” solutions of extremal problems involving graph homorphism densities. As mentioned in the introduction, such problems can be extremely difficult in full generality. Increasingly complex finitely forcible graphons are constructed in [GGKK15, GKK14, LS11], furthering the evidence in support of this claim. In particular, it is asked in [LS11] for which graphs F the infinite lexicographic power $\otimes^{\infty} F$ is finitely forcible. Our main technical result, Lemma 3.5, provides a partial answer to this question.

Extending the above definition, we say that a family \mathcal{W} of graphons is *finitely forcible* if there exists a quantum graph f such that $t(f; W') = 0$ for a graphon W' if and only if W' is weakly isomorphic to some graphon $W \in \mathcal{W}$. We say that such a quantum graph f *forces* \mathcal{W} .

3 Forcing lexicographic products

In this section we prove a series of technical results related to finite forcibility, which are used in the proof of our main result in the next section.

Let H and G be graphs. We say that a map $\phi : V(G) \rightarrow V(H)$ is a *folding of G into H* if for every pair of vertices $u, v \in V(G)$, either $\phi(u) = \phi(v)$, or $uv \in E(G)$ if and only if $\phi(u)\phi(v) \in E(H)$. That is ϕ preserves adjacencies and non-adjacencies except that it does not restrict adjacencies of the vertices of G mapped to the same vertex of H . Thus every strong homomorphism is a folding, but not vice versa. We say that a folding ϕ is *trivial* if it maps all the vertices to a single vertex.

A set $A \subseteq V(H)$ is called *homogeneous* in a graph H , if for every pair of distinct vertices $u, v \in A$, $N(u) \setminus A = N(v) \setminus A$, where $N(u)$ and $N(v)$ respectively denote the set of the neighbors of u and v . Equivalently, every vertex outside A is either adjacent to all the vertices in A , or has no neighbor in A . We call a graph H *prime*, if it does not contain any homogeneous sets A with $1 < |A| \leq |V(H)| - 1$. Note the following easy property of prime graphs.

Proposition 3.1. *Let $\phi : V(H) \rightarrow V(G)$ be a folding of a graph H into a graph G . Let A be a homogeneous set in G . Then $\phi^{-1}(A)$ is homogeneous in H . In particular, if H is prime, then either ϕ is trivial, or ϕ is injective, and thus is a strong homomorphism.*

We say that a graph H is *stringent* if H is prime, and furthermore, H does not have any non-identity automorphisms. If a graph H is stringent, then in particular the identity map is the only map from H to itself that preserves both adjacency and non-adjacency.

Proposition 3.2. *For any fixed $p \in (0, 1)$ the Erdős-Renyi random graph $G(n, p)$ is asymptotically almost surely stringent.*

Proof. Erdős and Renyi [ER63] have shown that the automorphism group of $G(n, p)$ is a.a.s. trivial. Thus it remains to show that $G(n, p)$ is a.a.s. prime.

If $G = G(n, p)$ is not prime then there exists a set S of size $2 \leq k \leq n - 1$ such that every vertex in $V(G) - S$ is either adjacent to all vertices in S or is not adjacent to any vertex in S . Thus the probability that G is not prime is at most

$$\sum_{k=2}^n \binom{n}{k} \left((1-p)^k + p^k \right)^{n-k} \leq \sum_{k=0}^n \binom{n}{k} \left((1-p)^2 + p^2 \right)^{n-k} = \left((1-p)^2 + p^2 \right)^n \xrightarrow{n \rightarrow \infty} 0,$$

as desired. □

We will need a statement similar to Proposition 3.1 involving maps from a prime graph into a graphon. Extending the definition of homogeneous sets in graphs, we say that a set $S \subseteq [0, 1]$ is *weakly homogeneous* in a graphon W , if there exists a function $e : [0, 1] \setminus S \rightarrow \{0, 1\}$ such that $W(x, y) = e(x)$ for almost every pair (x, y) with $x \in [0, 1] \setminus S$ and $y \in S$. We say that $S \subseteq [0, 1]$ is *strongly homogeneous* instead if $W(x, y) = e(x)$ holds for all pairs (x, y) as above. As we do not distinguish between weakly isomorphic graphons, and, in particular, between graphons which differ on a set of measure zero, it will be convenient for us to assume that a weakly homogeneous set is strongly homogeneous.

We are now ready to state a graphon analogue of Proposition 3.1.

Proposition 3.3. *Let W be a graphon, let S be a strongly homogeneous set in W , and let H be a fully labelled graph. If $t(H, \phi; W) > 0$ for an injective $\phi : V(H) \rightarrow [0, 1]$, then $\phi^{-1}(S)$ is homogeneous in H . In particular, if H is prime, then either $\phi^{-1}(S) = V(H)$, or $|\phi^{-1}(S)| \leq 1$.*

Note that if $(J_v)_{v \in V(F)}$ is an (F, μ) -partition of a graphon W , then, in particular, the set J_v is strongly homogeneous for every $v \in V(F)$. For a measure μ on the vertex set of a graph G and a positive integer k , define $\mathbf{m}_k(\mu) = \sum_{v \in V(G)} \mu^k(v)$. The preceding observation together with Proposition 3.3 implies the following.

Lemma 3.4. *Let (F, μ) be a weighted graph, let W be an (F, μ) -partitionable graphon, and let $(W_v)_{v \in V(F)}$ be the collection of bags of some (F, μ) -partition of W . Then*

$$t(H; W) = t(H; F, \mu) + \sum_{v \in V(F)} \mu^k(v) t(H; W_v) \quad (3)$$

for any prime graph H on k vertices.

In particular, if $W = \otimes_{i=1}^{\infty} (F_i, \mu_i)$ for some sequence of weighted graphs $(F_i, \mu_i)_{i \in \mathbb{N}}$, then

$$t(H; W) = \sum_{i=1}^{\infty} \left(\prod_{j=1}^{i-1} \mathbf{m}_k(\mu_j) \right) t(H; F_i, \mu_i). \quad (4)$$

We say that a finite collection \mathcal{F} of graphs is *n-stringent* if each graph $F \in \mathcal{F}$ is stringent, and $V(F) = [n]$ for every $F \in \mathcal{F}$. Let $\mu : [n] \rightarrow (0, 1)$ be a probabilistic measure. We define $\text{Lex}(\mathcal{F}, \mu)$ to be the set of all infinite lexicographic products $\otimes_{i=1}^{\infty} J_i$, where for every $i \geq 1$ we have $J_i = (F, \mu)$ for some $F \in \mathcal{F}$. The following lemma is the main result of this section.

Lemma 3.5. *The family $\text{Lex}(\mathcal{F}, \mu)$ is finitely forcible for every n-stringent collection of graphs \mathcal{F} and every probabilistic measure $\mu : [n] \rightarrow (0, 1)$.*

The remainder of the section is occupied with the proof of Lemma 3.5. The proof relies, in particular, on the following simple observation about finitely forcible families.

Proposition 3.6. *Let $\mathcal{W}_1, \mathcal{W}_2$ be finitely forcible families of graphons closed under weak isomorphism. Then $\mathcal{W}_1 \cap \mathcal{W}_2$ is finitely forcible.*

Proof. If f_i forces \mathcal{W}_i for $i = 1, 2$, then $f_1^2 + f_2^2$ forces $\mathcal{W}_1 \cap \mathcal{W}_2$. □

A graphon is called *random-free* if it is $\{0, 1\}$ -valued almost everywhere.

Proposition 3.7. *Let \mathcal{F} be a family of n-stringent graphs. Then $\text{Lex}(\mathcal{F}, \mu) \subseteq \mathcal{W}_{\text{rf}}$ for a finitely forcible family of random-free graphons \mathcal{W}_{rf} .*

Proof. Consider a bipartite graph H with a bipartition $V(H) = X \cup Y$, and let the quantum graph \overline{H} be the sum of all graphs that can be obtained from H by adding only edges whose end-points are either both in X or both in Y . Note that $t(\overline{H}; G)$ is the probability that a random map $\phi : V(H) \rightarrow V(G)$ satisfies $\phi(x)\phi(y) \in E(G)$ if and only if $xy \in E(H)$ for all pairs $x \in X, y \in Y$. It is shown in [LS10] that $t(\overline{H}; W) = 0$ implies that W is random-free. Thus our task reduces to finding a bipartite graph H such that $t(\overline{H}; W) = 0$ for all $W \in \text{Lex}(\mathcal{F}, \mu)$.

It is well-known that with high probability the size of the largest independent set in the Erdős-Rényi random graph $G(m, 1/2)$ is bounded by $2 \log m$. Thus by Proposition 3.2, there exists a stringent graph K with $m := |V(K)| > n$ whose independence number is smaller than $2 \log m$. Let H be the bipartite graph obtained from K by splitting every vertex i of K into two vertices u_i and v_i , and connecting u_i to v_j if and only if $ij \in E(K)$. The sets $X = \{u_i : i \in V(K)\}$ and $Y = \{v_i : i \in V(K)\}$ form a bipartition of H . Suppose $t(\overline{H}; W) > 0$ for some $W \in \text{Lex}(\mathcal{F}, \mu)$. Then for some $F \in \mathcal{F}$, there exists a map $\phi : V(H) \rightarrow V(F)$ such that $t(\overline{H}, \phi; F, \mu) > 0$, and ϕ does not map all the vertices of H to a single vertex. It follows from $t(\overline{H}, \phi; F, \mu) > 0$ that for all $u_i \in X$ and $v_j \in Y$ satisfying $\phi(u_i) \neq \phi(v_j)$, we have $\phi(u_i)\phi(v_j) \in E(F)$ if and only if $u_i v_j \in E(H)$. Pick $a \in V(F)$ such that $2m/n \leq |\phi^{-1}(a)| < 2m$. Note that $\{i : \phi(u_i) = \phi(v_i) = a\}$ is a homogeneous set in K , and thus is of size at most 1. But both sets $\{i : \phi(u_i) = a, \phi(v_i) \neq a\}$ and $\{i : \phi(u_i) \neq a, \phi(v_i) = a\}$ are independent sets in K , and thus they are of size at most $2 \log m$. For sufficiently large m , this contradicts the choice of a . □

By Proposition 3.6, to prove Lemma 3.5 it suffices to construct finitely forcible families of graphons $\mathcal{W}_1, \dots, \mathcal{W}_l$ for some l such that $\text{Lex}(\mathcal{F}, \mu) = \cap_{i=1}^l \mathcal{W}_i$. Defining these families will require a technical definition.

Definition 3.8. Let k be a positive integer, (F, μ) a weighted graph, and W a graphon. We say that W is an (F, μ, k) -split graphon if for (almost) every map $\phi : V(F) \rightarrow (0, 1)$ with $t(F, \phi; W) > 0$, there exists a homogeneous set S , and a partition $(J_v)_{v \in V(F)}$ of S such that the following holds:

- (S1) $(J_v)_{v \in V(F)}$ corresponds to an (F, μ) -partition of $W[S]$;
- (S2) $\phi(v) \in J_v$ for every $v \in V(F)$;
- (S3) $t(H; W[J_v]) = t(H; W[J_u])$ for every pair of vertices $u, v \in V(F)$, and every graph H on k vertices.

Let $\mathcal{S}_k(F, \mu)$ denote the family of all (F, μ, k) -split graphons.

Lemma 3.9. Let k be a positive integer, and let F be a stringent graph. Let \mathcal{W}_{rf} be a finitely forcible family of random-free graphons. Then $\mathcal{S}_k(F, \mu) \cap \mathcal{W}_{\text{rf}}$ is finitely forcible for every probabilistic measure $\mu : V(F) \rightarrow (0, 1)$.

Proof. By Proposition 3.7 it suffices to show the existence of a quantum graph f such that a random-free graphon W satisfies $t(f; W) = 0$ if and only if $W \in \mathcal{S}_k(F, \mu)$. Assume $V(F) = [n]$. We will think of F as an $[n]$ -labelled graph.

Consider a map $\phi : [n] \rightarrow [0, 1]$ and a $\{0, 1\}$ -valued graphon W such that $t(F, \phi; W) = 1$. For every vertex i of F , let J_i be the set of points $x \in [0, 1]$ such that $W(x, \phi(j)) = W(\phi(i), \phi(j))$ for all $j \in [n]$. Let J_\circ be the set of $x \in [0, 1]$ such that $W(x, \phi(j)) = 0$ for all $j \in [n]$, and let J_\bullet be the set of $x \in [0, 1]$ such that $W(x, \phi(j)) = 1$ for all $j \in [n]$. Since F is stringent, the sets $J_1, \dots, J_n, J_\circ, J_\bullet$ are disjoint. Let $S := J_1 \cup \dots \cup J_n$. We would like to force the sets S and J_1, \dots, J_n to satisfy the conditions of Definition 3.8. We will introduce a few conditions of the form $t(f_i, \phi; W) = a_i$ that would imply this. Consequently by the discussion in the last paragraph of Section 2.2, setting $f = \llbracket \sum (f_i - a_i)^2 \rrbracket$ will yield the desired result.

Define $h = \sum_{1 \leq i}^n \tilde{F}^{\{i\}}$, let F° be obtained from F by adding an unlabeled isolated vertex, and let F^\bullet be obtained from F by connecting a new unlabeled vertex to all the vertices of F . For every graph H on k vertices, and every $i \in [n]$, let $F \oplus_i H$ be the $[n]$ -partially labelled graph that is obtained from F by first creating k unlabeled twins for the vertex i , and then implanting a copy of H on these unlabeled vertices. Let $F \tilde{\oplus}_i H$ be sum of all the 2^k graphs that are obtained from $F \oplus_i H$ by possibly adding some edges between the vertex i and the unlabeled vertices. Note the following

1. The condition $t(h + F^\circ + F^\bullet, \phi; W) = 1$ implies that $J_1 \cup \dots \cup J_n \cup J_\circ \cup J_\bullet = [0, 1]$ up to a measure zero difference.
2. The condition $t(\tilde{F}^{\{i\}} - \mu(i)h, \phi; W) = 0$ implies that $\lambda(J_i) = \mu(i)\lambda(J_1 \cup \dots \cup J_n)$, and thus J_i are of desired measures.
3. The condition $t(\tilde{F}^{\{i,j\}} - \tilde{F}^{\{i\}}\tilde{F}^{\{j\}}, \phi; W) = 0$ implies that for almost every $x \in J_i$ and $y \in J_j$, we have $W(x, y) = 1$ if and only if $ij \in E(F)$. This together with the previous condition imply that $(J_i)_{i \in V(F)}$ is an (F, μ) -partition of $W[S]$.
4. It can be similarly forced that $W(x, y) = 1$ for all $x \in J_i$ and $y \in J_\bullet$, and $W(x, y) = 0$ for all $x \in J_i$ and $y \in J_\circ$. Thus S is a homogeneous set.
5. The condition $\mu(i)^{-k}t(F \tilde{\oplus}_i H, \phi; W) = \mu(j)^{-k}t(F \tilde{\oplus}_j H, \phi; W)$ implies $t(H; W[J_i]) = t(H; W[J_j])$, and thus (S3) can be forced in this manner.

As we mentioned above, the desired quantum graph can be constructed from these conditions. \square

Before finally stating the proof of Lemma 3.5, let us make a simple observations regarding (F, μ, k) -split graphons.

Lemma 3.10. *Let F_1 and F_2 be stringent graphs, and let W be an (F_1, μ, k) -split graphon. Let S , and $(J_v)_{v \in V(F_1)}$ be as in Definition 3.8 for some $\phi : V(F_1) \rightarrow [0, 1]$. Then the following hold.*

- (i) *If W is an (F_2, μ, k) -split graphon, then so is $W[J_v]$.*
- (ii) *If $\phi : V(F_2^{(d)}) \rightarrow [0, 1]$ satisfies $t(\tilde{F}_2^{(d)}, \phi; W) > 0$, then either $\phi(V(F_2^{(d)})) \subseteq S$, or there is a vertex $v \in V(F_2)$ such that for all the vertices $u \in V(F_2^{(d)})$ that are not copies of v , we have $\phi(u) \notin S$.*

Proof. To verify (i) pick a vertex $v_0 \in V(F_1)$, and consider a map ψ with $t(F_2, \psi; W[J_{v_0}]) > 0$ and its corresponding map $\psi' : V(F_i) \rightarrow J_{v_0}$ with $t(F_2, \psi'; W) > 0$. If W is an (F_2, μ, k) -graphon, one can find a homogeneous \tilde{S} in W and a partition $(\tilde{J}_v)_{v \in V(F)}$ of \tilde{S} satisfying (S1), (S2), and (S3). Since F_2 is stringent and J_{v_0} is a homogeneous set for W , by Proposition 3.3 we must have $\tilde{J}_v \subseteq J_{v_0}$ for all v . Thus $W[J_{v_0}]$ is an (F_2, μ, k) -split graphon.

To verify (ii), consider a map $\phi : V(F_2^{(d)}) \rightarrow [0, 1]$ satisfying $t(\tilde{F}_2^{(d)}, \phi; W) > 0$. Consider a copy of F_2 in $F_2^{(d)}$. If ϕ maps two different vertices from this copy to points in S , then since S is homogeneous and F_2 is stringent, it must map all the vertices of this copy into S . This in turn implies that all the vertices of $F_2^{(d)}$ must be mapped into S . \square

Proof of Lemma 3.5. Let $\mathcal{F} = (F_i)_{i=1}^l$ be the family of n -stringent graphs; that is for every i , the graph F_i is stringent and $V(F_i) = [n]$. We consider the graphs F_1, \dots, F_n as fully labeled graphs. Let d be a positive integer chosen to be sufficiently large to satisfy the inequalities appearing later in the proof. By Lemma 3.9, the class of graphons $\mathcal{W}^* = \cap_{i=1}^l \mathcal{S}_{nd}(F_i, \mu) \cap \mathcal{W}_{\text{rf}}$ is finitely forcible, where \mathcal{W}_{rf} is as in Proposition 3.7. Clearly $\text{Lex}(\mathcal{F}, \mu) \subseteq \mathcal{W}^*$. Let f_1 be a quantum graph that forces \mathcal{W}^* . We will show that there exists a quantum graph f_2 such that $t(f_2; W) = 0$ for $W \in \mathcal{W}^*$ if and only if $W \in \text{Lex}(\mathcal{F}, \mu)$. This will imply that $f_1^2 + f_2^2$ forces $\text{Lex}(\mathcal{F}, \mu)$, and thus the lemma.

Let $f_2' = \llbracket \sum_{i=1}^l \tilde{F}_i^{(d)} \rrbracket$, and set $\rho = \prod_{i=1}^n \mu(i)$ and $m = \frac{\rho^d}{1 - \mathbf{m}_{nd}(\mu)}$. By (4) we have $t(f_2'; W) = m$ for every $W \in \text{Lex}(\mathcal{F}, \mu)$. We will show that, if $t(f_2'; W) \geq m$ for some $W \in \mathcal{W}^*$, then $W \in \mathcal{W}^*$. This will imply that the quantum graph $f_2 = f_2' - m$ is as desired.

Let $m^* = \max_{W \in \mathcal{W}^*} t(f_2'; W)$. Our first goal is to show that $m^* = m$. Consider $W \in \mathcal{W}^*$ with $t(f_2'; W) = m^*$, and pick $i \in [l]$ so that $t(\llbracket \tilde{F}_i^{(d)} \rrbracket; W) \geq m/l$. Consequently, there exists $\phi : [n] \rightarrow [0, 1]$ with $t(\tilde{F}_i^{(d)}, \phi; W) \geq m/l$; in particular, ϕ satisfies the conditions in the definition of the (F_i, μ, nd) -split graphon. Consider the homogeneous set S , and the partition $(J_v)_{v \in [n]}$ as in Definition 3.8. Note that $W[S] \in \mathcal{W}^*$, by construction, and moreover since F_1, \dots, F_l are stringent graphs, by Lemma 3.10 (i) for every $v \in [n]$, $W[J_v] \in \mathcal{W}^*$.

Set $\varepsilon = 1 - \lambda(S)$, and note that

$$\frac{m}{l} \leq t(\tilde{F}_i^{(d)}, \phi; W) \leq (1 - \varepsilon)^{n(d-1)}.$$

Therefore,

$$\varepsilon \leq 1 - \left(\frac{m}{l}\right)^{\frac{1}{n(d-1)}} = 1 - \left(\frac{\rho^d}{l(1 - \mathbf{m}_{nd}(\mu))}\right)^{\frac{1}{n(d-1)}} \leq \frac{1}{2}, \quad (5)$$

provided that d is sufficiently large. However, by Lemma 3.10 (ii), we have

$$\begin{aligned} m^* &= t(f_2'; W) \leq (1 - \varepsilon)^{nd} t(f_2'; W[S]) + l(n+1)\varepsilon^{n(d-1)} \\ &\leq (1 - \varepsilon)^{nd} m^* + l(n+1)\varepsilon^{n(d-1)} \leq (1 - \varepsilon)m^* + l(n+1)\varepsilon^{d-1}. \end{aligned} \quad (6)$$

If $\varepsilon > 0$, then (5) and (6) imply that

$$m \leq m^* \leq l(n+1)\varepsilon^{d-2} \leq \frac{l(n+1)}{2^{d-2}};$$

a contradiction for sufficiently large d . Hence we conclude that $\varepsilon = 0$. Thus W is F_i -partitionable, and by Lemma 3.4,

$$m^* = t(f'_2; W) = \rho^d + \sum_{v \in [n]} \mu(v)^{nd} t(f'_2; W[J_v]) \leq \rho^d + \sum_{v \in [n]} \mu(v)^{nd} m^* = \rho^d + \mathbf{m}_{nd}(\mu) m^*. \quad (7)$$

It follows from the definition of m that $m^* = m$, and that all the inequalities in (7) hold with equality; that is $t(f'_2; W[J_v]) = m$ for all $v \in [n]$. Thus, applying the preceding argument to $W[J_v]$ for every $v \in [n]$, we conclude that each $W[J_v]$ is $(F_{j(v)}, \mu)$ -partitionable for some $1 \leq j(v) \leq l$.

Note that $t(\llbracket \tilde{F}_{j(v)}^{(d)} \rrbracket; W[J_v]) \geq \rho^d$. Moreover, by Lemma 3.4, we obtain that if $k \neq j(v)$, then

$$t(\llbracket \tilde{F}_k^{(d)} \rrbracket; W[J_v]) \leq \sum_{u \in [n]} \mu(u)^{nd} m = \mathbf{m}_{nd}(\mu) m < \rho^d,$$

as long as d is large enough so that $\mathbf{m}_{nd}(\mu) < \frac{1}{2}$. Thus by the property (S3) of (F_i, μ, nd) -split graphons we deduce that there exists j such that $W[J_v]$ is F_j -partitionable for every $v \in [n]$. Repeating the same argument recursively we deduce $W \in \text{Lex}(\mathcal{F}, \mu)$, as desired. \square

4 The main result

We are now ready to state and prove our main result, which was informally described in the introduction, using the tools developed in Section 3.

Theorem 4.1. *There exists a quantum graph f such that the function $\psi: [0, 1] \rightarrow \mathbb{R}$ defined as $\psi(x) = \min t(K_2; W)$ where the minimum is over all graphons W with $t(f; W) = 0$ and $t(P_4; W) = x$ is nowhere differentiable in the interval $[\alpha_1, \alpha_2]$ for some $0 \leq \alpha_1 < \alpha_2 \leq 1$.*

Proof. Let F_1 and F_2 be two stringent graphs with $V(F_1) = V(F_2) = [n]$. Let μ be the unique probability measure on $[n]$ satisfying $\sum_{i=1}^n \mu(i)^4 = \frac{1}{2}$ and $\mu(2) = \dots = \mu(n)$. Let $\alpha_i = t(P_4; F_i, \mu)$, and $\beta_i = \frac{t(P_2; F_i, \mu)}{1 - \mathbf{m}_2(\mu)}$. We assume $\alpha_1 < \alpha_2$, and $\beta_1 < \beta_2$. Indeed if one takes F_1 to be a random graph $G(n, 1/3)$, and F_2 to be a random graph sampled $G(n, 1/2)$, then it is straightforward to see, using Proposition 3.2 that F_1 and F_2 satisfy the above properties asymptotically almost surely.

Let $\mathcal{F} = \{(F_1, \mu), (F_2, \mu)\}$. Then \mathcal{F} is stringent, and thus by Lemma 3.5 there exists a quantum graph f such that $t(f; W') = 0$ for a graphon W' if and only if W' is weakly isomorphic to some graphon $W \in \text{Lex}(\mathcal{F})$. We will show that f, α_1, α_2 satisfy the theorem. By the choice of f , we have $\psi(x) = \min t(K_2; W)$, where the minimum is over all $W \in \text{Lex}(\mathcal{F})$ with $t(P_4; W) = x$.

Consider $W = \otimes_{i=1}^\infty J_i$, where for every $i \geq 1$ either $J_i = (F_1, \mu)$ or $J_i = (F_2, \mu)$. Let $T = \{i \in \mathbb{N} : J_i = (F_1, \mu)\}$. Let $\lambda = \sum_{i \in T} \frac{1}{2^{i-1}}$, and let $\gamma = \mathbf{m}_2(\mu) > 1/2$. By Lemma 3.4 we have

$$t(P_4; W) = \sum_{j=1}^\infty \frac{1}{2^{j-1}} t(P_4; J_j) = \sum_{i \in T} \frac{1}{2^{i-1}} \alpha_1 + \sum_{i \in \mathbb{N}-T} \frac{1}{2^{i-1}} \alpha_2 = \lambda \alpha_1 + (1 - \lambda) \alpha_2, \quad (8)$$

and, similarly,

$$t(K_2; W) = \sum_{j=1}^\infty \gamma^{j-1} t(K_2; J_j) = \beta_2 + (\beta_1 - \beta_2)(1 - \gamma) \sum_{i \in T} \gamma^{i-1}.$$

The identity (8) shows that when λ is not a dyadic rational, and $x = \lambda\alpha_1 + (1 - \lambda)\alpha_2$ a convex combination of α_1 and α_2 , then there exists a unique $W \in \text{Lex}(\mathcal{F})$ that satisfies $t(P_4, W) = x$, and if λ is dyadic then there are two such $W \in \text{Lex}(\mathcal{F})$ corresponding to the two different binary expansions of λ . Out of the two choices of $(\lambda_i)_{i \in \mathbb{N}}$, the one that corresponds to a finite T minimizes $t(K_2; W)$.

Now it can be easily seen now that for every $x \in [\alpha_1, \alpha_2]$,

$$\limsup_{h \rightarrow 0} \frac{\psi(x + h) - \psi(x)}{h} = +\infty,$$

finishing the proof of the theorem. □

5 Concluding remarks

Differentiability of the boundary of $T(\mathcal{F})$. In the main theorem of this article, Theorem 4.1, we showed that there exists a quantum graph f such that the set of points $(t(K_2; W), t(P_4; W))$ for all graphons W that satisfy $t(f; W) = 0$ is the union of a curve that has a nowhere differentiable part and a countable set of points. The region defined by f is of measure 0, however it is not difficult to slightly modify f to obtain a region of positive measure whose boundary is nowhere differentiable in some parts.

Let $f = \sum_{F \in \mathcal{F}} \lambda_F F$ be the quantum graph in Theorem 4.1. The theorem says that the projection of the intersection of $T(\mathcal{F} \cup \{K_2, P_4\})$ with the hyperplane defined by $\sum_{F \in \mathcal{F}} \lambda_F x_F = 0$ to the coordinates (x_{K_2}, x_{P_4}) is a curve that is nowhere differentiable in certain parts plus a countable set of points. Unfortunately this does not imply anything about differentiability of the boundary of $T(\mathcal{F} \cup \{K_2, P_4\})$. The main question still remain open: For a finite collection \mathcal{F} of graphs, is the boundary of the region $T(\mathcal{F})$ almost everywhere differentiable?

Finite forcibility of lexicographic powers. Lemma 3.5 implies that $\otimes^\infty F$ is finitely forcible for every stringent graph F , giving a partial answer to a question of Lovász and Szegedy [LS11] mentioned in Section 2.5. We believe that one can extend the argument to show that $\otimes^\infty F$ is finitely forcible for every prime F . Dealing with homogeneous sets, however, presents a major technical issue.

Acknowledgement. The second author thanks Liana Yepremyan for discussions related to the subject of this paper.

References

- [ELS79] Paul Erdős, László Lovász, and Joel Spencer, *Strong independence of graphcopy functions*, Graph theory and related topics (Proc. Conf., Univ. Waterloo, Waterloo, Ont., 1977), Academic Press, New York-London, 1979, pp. 165–172. MR 538044 (81b:05060) [1](#)
- [ER63] P. Erdős and A. Rényi, *Asymmetric graphs*, Acta Math. Acad. Sci. Hungar **14** (1963), 295–315. MR 0156334 [5](#)
- [GGH⁺16] Roman Glebov, Andrzej Grzesik, Ping Hu, Tamas Hubai, Daniel Kral, and Jan Volec, *Densities of 3-vertex graphs*, arXiv:1610.02446. [1](#)
- [GGKK15] Roman Glebov, Andrzej Grzesik, Tereza Klimošová, and Daniel Král', *Finitely forcible graphons and permutons*, J. Combin. Theory Ser. B **110** (2015), 112–135. MR 3279390 [4](#)
- [GKK14] Roman Glebov, Tereza Klimošová, and Daniel Král', *Infinite dimensional finitely forcible graphon*, arXiv:1404.2743. [4](#)
- [HLN⁺14] Hao Huang, Nati Linial, Humberto Naves, Yuval Peled, and Benny Sudakov, *On the 3-local profiles of graphs*, J. Graph Theory **76** (2014), no. 3, 236–248. MR 3200287 [1](#)

- [HN11] Hamed Hatami and Serguei Norine, *Undecidability of linear inequalities in graph homomorphism densities*, J. Amer. Math. Soc. **24** (2011), no. 2, 547–565. MR 2748400 [1](#)
- [Lov12] László Lovász, *Large networks and graph limits*, American Mathematical Society Colloquium Publications, vol. 60, American Mathematical Society, Providence, RI, 2012. MR 3012035 [1](#)
- [LS06] László Lovász and Balázs Szegedy, *Limits of dense graph sequences*, J. Combin. Theory Ser. B **96** (2006), no. 6, 933–957. MR MR2274085 (2007m:05132) [2](#)
- [LS10] ———, *Regularity partitions and the topology of graphons*, An irregular mind, Bolyai Soc. Math. Stud., vol. 21, János Bolyai Math. Soc., Budapest, 2010, pp. 415–446. MR 2815610 [6](#)
- [LS11] L. Lovász and B. Szegedy, *Finitely forcible graphons*, J. Combin. Theory Ser. B **101** (2011), no. 5, 269–301. MR 2802882 [4](#), [10](#)
- [Raz07] Alexander A. Razborov, *Flag algebras*, J. Symbolic Logic **72** (2007), no. 4, 1239–1282. MR MR2371204 (2008j:03040) [1](#)
- [Raz08] ———, *On the minimal density of triangles in graphs*, Combin. Probab. Comput. **17** (2008), no. 4, 603–618. MR MR2433944 (2009i:05118) [1](#)
- [Rei16] Christian Reiher, *The clique density theorem*, Ann. of Math. (2) **184** (2016), no. 3, 683–707. MR 3549620 [1](#)